



US 20030115424A1

(19) **United States**(12) **Patent Application Publication****Bachand et al.**(10) **Pub. No.: US 2003/0115424 A1**(43) **Pub. Date: Jun. 19, 2003**(54) **SNOOP BLOCKING FOR CACHE COHERENCY**(76) **Inventors: Derek T. Bachand, Portland, OR (US);  
Paul Breuder, Beaverton, OR (US);  
Matthew A. Fisch, Portland, OR (US)**

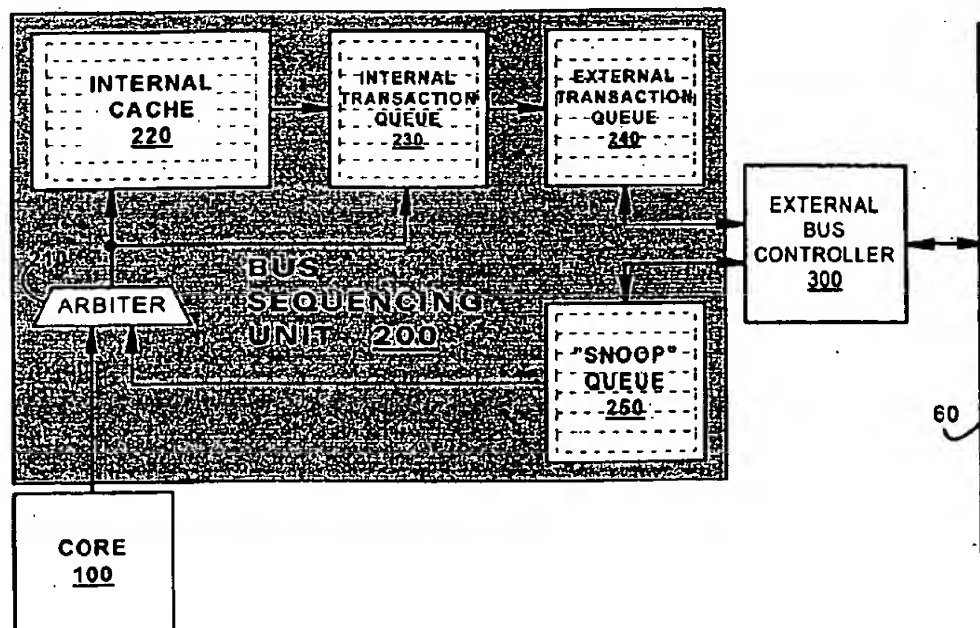
No. 09/210,641, filed on Dec. 14, 1998, now Pat. No. 6,460,119, which is a continuation-in-part of application No. 08/999,242, filed on Dec. 29, 1997, now Pat. No. 6,209,068.

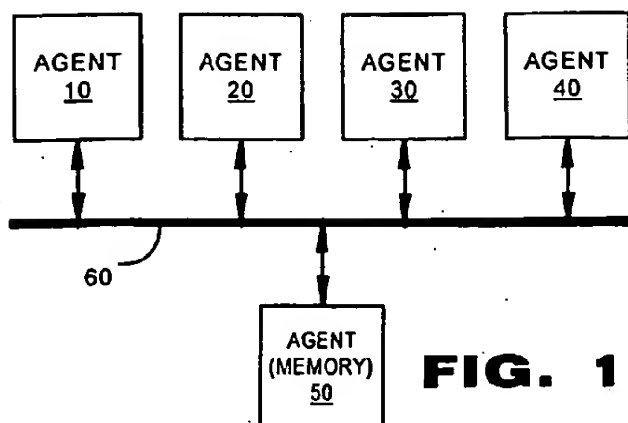
**Publication Classification**(51) **Int. Cl.<sup>7</sup> ..... G06F 12/00; G06F 12/14**(52) **U.S. Cl. .... 711/146; 711/163**

Correspondence Address:

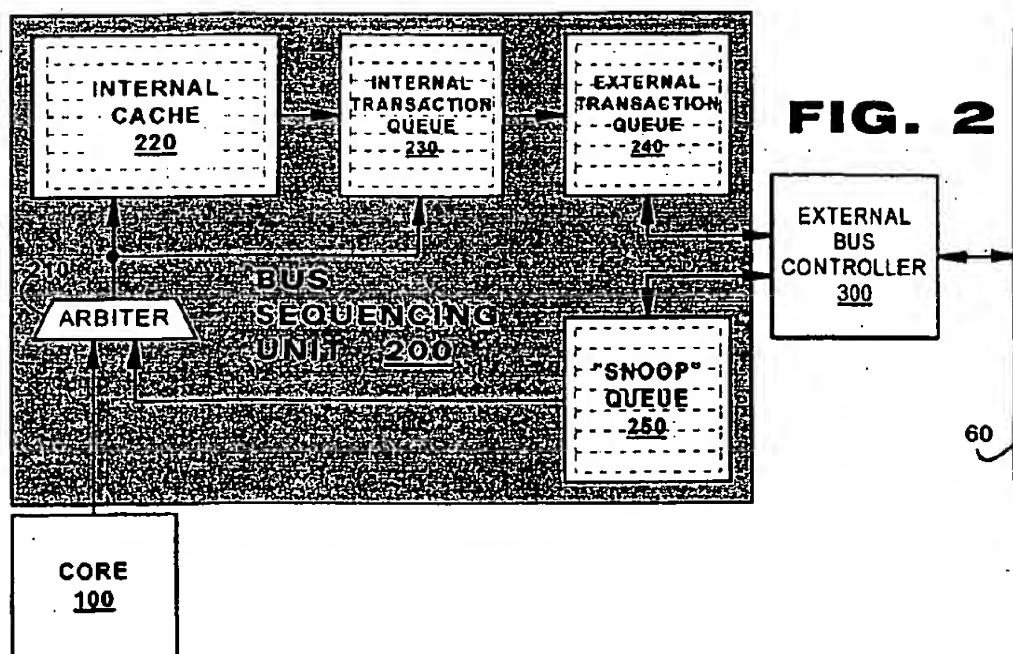
**KENYON & KENYON****1500 K STREET, N.W., SUITE 700****WASHINGTON, DC 20005 (US)**(21) **Appl. No.: 10/352,946**(22) **Filed: Jan. 29, 2003****Related U.S. Application Data**(63) **Continuation of application No. 10/215,599, filed on Aug. 9, 2002, which is a continuation of application**(57) **ABSTRACT**

In a method of processing a bus transaction, an address is retrieved from the bus transaction and referred to a queue of pending transaction. A match indicator signal is returned from the queue. If the match indicator signal indicates a match, a snoop probe for the bus transaction is blocked.

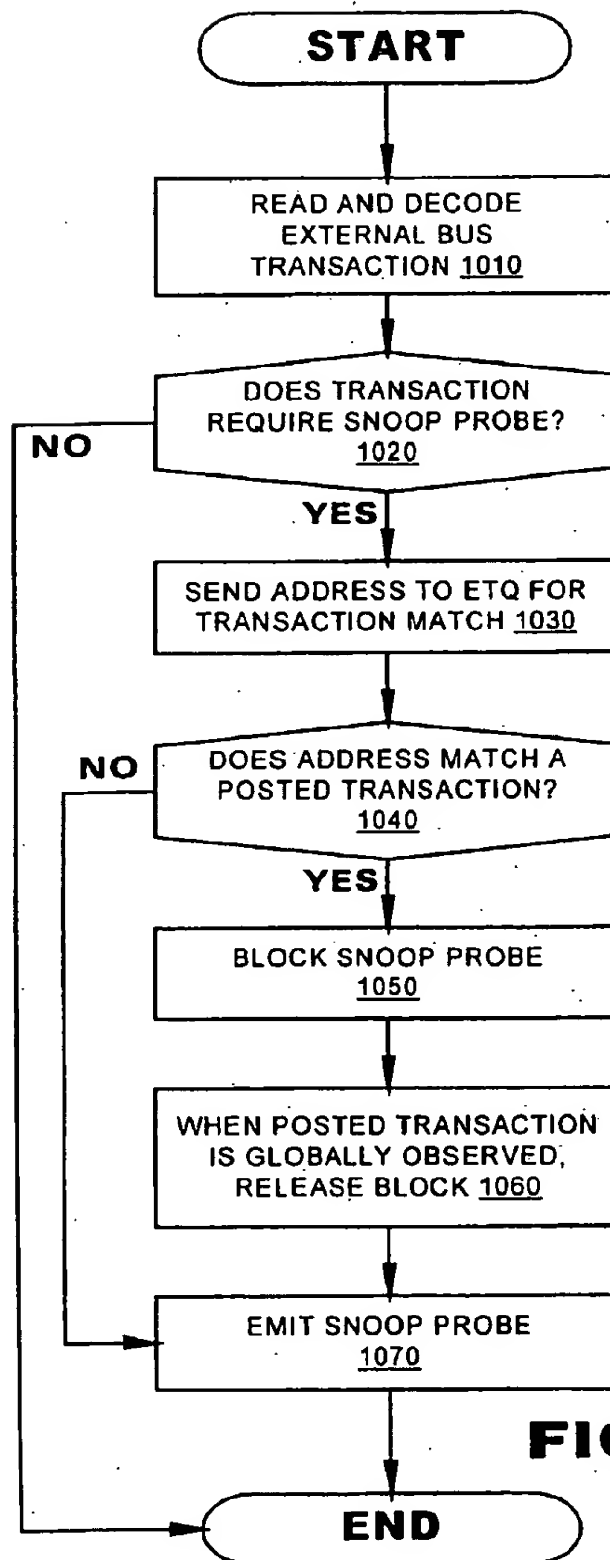




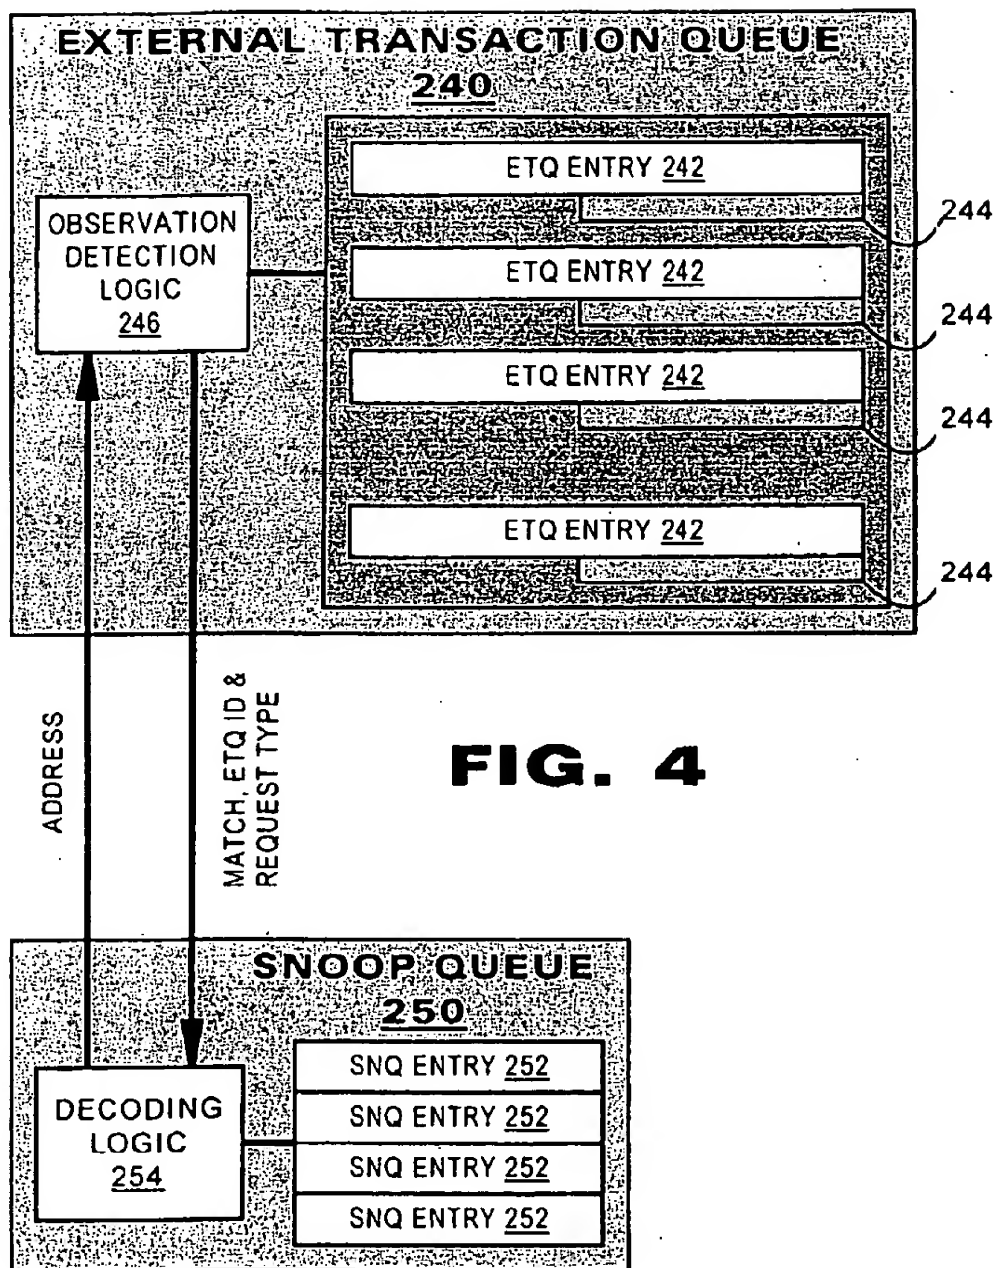
**FIG. 1**



**FIG. 2**



**FIG. 3**



## SNOOP BLOCKING FOR CACHE COHERENCY

## BACKGROUND

[0001] The present invention relates to a cache coherency technique in an agent using a pipelined bus.

[0002] As is known, many modern computing system employ a multi-agent architecture. A typical system is shown in FIG. 1. There, a plurality of agents 10-50 communicate over an external bus 60 according to a predetermined bus protocol. "Agents" may include general purpose processors, chipsets for memory and/or input output devices or other integrated circuits that process data requests. The bus 60 may be a "pipelined" bus in which several transactions may be in progress at once. Each transaction progresses through a plurality of stages but no two transactions are in the same stage at the same time. The transactions complete in order. With some exceptions, transactions generally do not "pass" one another as they progress on the external bus 60.

[0003] In a multiple-agent system, two or more agents may have need for data at the same memory location at the same time. The agents 10-50 operate according to cache coherency rules to ensure that each agent 10 uses the most current copy of the data available to the system. According to many cache coherency systems, each time an agent 10 stores a copy of data, it assigns to the copy a state indicating the agent's rights to read and/or modify the data.

[0004] For example, the Pentium® Pro processor, commercially available from Intel Corporation, operates according to the "MESI" cache coherency scheme. Each copy of data stored in an agent 10 is assigned one of four states including:

[0005] Invalid—Although an agent 10 may have cached a copy of the data, the copy is unavailable to the agent. The agent 10 may neither read nor modify an invalid copy of data.

[0006] Shared—The agent 10 stores a copy of data that is valid and possesses the same value as is stored in external memory. An agent 10 may only read data in shared state. Copies of the data may be stored with other agents also in shared state. An agent 10 may not modify data in shared state without first performing an external bus transaction to gain exclusive ownership of the data.

[0008] Modified—The agent 10 stores a copy of data that is valid and "dirty." A copy cached by the agent 10 is more current than the copy stored in external memory. When an agent 10 stores data in modified state, no other agents possess a valid copy of the data.

[0009] Agents 10-50 exchange cache coherency messages, called "snoop responses," during external bus transactions. The snoop responses identify whether other agents possess copies of requested data and, if so, the states in which the other copies are held. For example, when an agent 10 requests data held in modified state by another agent 20, the other agent 20 may provide the data to the requesting agent in an implicit writeback. Ordinarily, data is provided to requesting agents 10 by the external memory 50. The modified data is the most current copy of data available to the system and should be transferred to the requesting agent 10 in response to a data request.

[0010] When external bus transactions cause an agent to change the state assigned to a copy of data, state changes occur after snoop responses are globally observed.

[0011] As an example, consider a "read for ownership" request issued by an agent 10. Initially, an agent 10 may store the requested data in an invalid state. The agent 10 has a need for the data and issues a bus transaction requesting it. The agent 10 receives snoop responses from other agents 20-40. When the snoop responses are received, the transaction is globally observed. The agent 10 marks the requested data as held in exclusive state. The agent 10 may mark the data even though it has not yet received the requested data. For example, in known processors, data is transferred in a data phase of a transaction following a snoop phase. Before the data is received, an entry of an internal cache (not shown) is reserved for the data. A state field in the external transaction queue is marked as exclusive when the transaction is globally observed and before the requested data is 10 received, but the state field in the reserved cache entry is not marked exclusive until the data is filled into the cache.

[0012] Certain boundary conditions arise when state transitions are triggered by the receipt of snoop responses. An example is shown in the following table using the Pentium® Pro bus protocol:

		Bus Cycles										
		1	2	3	4	5	6	7	8	9	10	11
Transaction No. 1	Req	Req	Err		Snoop	Stall		Snp	Rcsp	Data	X	
State in Agent 10	I	I	I	I	I	I	I	E	E	E	E	E
Transaction No. 2	X	X	Req	Req	Err	Snoop	Stall		Snp	Resp	Data	
State in Agent 20	I	I	I	I	I	I	I	I	E	E	E	E

[0007] Exclusive—The agent 10 stores a copy of data that is valid and may possess the same value as is stored in external memory. When an agent 10 caches data in exclusive state, it may read and modify the data without an external cache coherency check.

[0013] In the boundary condition, without some sort of preventative measure, two different agents 10 and 20 in the system could mark a copy of the same data in exclusive state. To do so would violate cache coherency. Assume that two agents 10 and 20 post read requests to a single piece of

data. The first agent 10 posts the request as explained above. When the first transaction concludes its request phase, the second agent 20 posts a second transaction for the same data.

[0014] Assume further that the snoop phase of the first transaction is stalled by a snoop stall. A snoop stall signal occurs when an agent (say, agent 30) requires additional time to generate snoop results. Although the first agent 10 may reserve a cache entry for the requested data, the agent 10 does not mark the requested data as exclusive until snoop results for its transaction are received. When snoop results eventually are received for the first transaction (in clock 8), the first agent 10 will mark the data as held in exclusive state. However, the first agent 10 observes the second transaction in clock 3. If it performs internal snoop inquiries for the second transaction before the first transaction is globally observed, its snoop response would indicate that it does not possess a valid copy of the data. The second agent 20 also could mark the data as exclusive. Having two agents 10, 20 each store data in exclusive state violates the MESI cache coherency rules because each agent 10, 20 could modify its copy of the data without notifying the other via a bus transaction.

[0015] The coherency violation can arise if an agent 10 begins internal snoop inquiries before its previous transaction to the data is globally observed. Thus, the error can be avoided if the snoop inquiries related to the second transaction are blocked until a prior conflicting transaction related to the same data is globally observed.

[0016] The Pentium® Pro processor includes a snoop queue to manage cache coherency and generate snoop responses. The snoop queue buffers all transactions posted on the external bus. For new transactions, the snoop queue compares the address of the new transaction to addresses of transactions that it previously stored to determine whether the addresses match. If so, and if the previous transaction were not globally observed, the snoop queue blocks a snoop probe for the new transaction. The block remains until snoop results for the prior pending transaction are received.

[0017] The Pentium® Pro processor's snoop queue is large. The snoop queue possesses a queue entry for as many transactions as can be pending simultaneously on the external bus. It consumes a large area when the Pentium® Pro processor is manufactured as an integrated circuit. In future processors, it will be desirable to increase the pipeline depth of the external bus to increase the number of transactions that may proceed simultaneously thereon. However, increasing the depth of the external bus becomes expensive if it also requires increasing the depth of the snoop queue.

[0018] The Pentium® Pro processor's snoop queue fills quickly during operation. The snoop queue buffers not only requests from other agents but also requests posted by the agent to which the snoop queue belongs. Because the Pentium® Pro includes an external transaction queue that monitors transactions issued by the processor, the snoop queue's design is considered suboptimal.

[0019] Accordingly, the inventors perceived a need in the art for a snoop queue in an agent that possesses a depth that is independent of the pipeline depth of the agent's external bus. There is a need in the art for such a snoop queue, however, that maintains cache coherency and insures that, when two bus transactions related to the same address are

pending on the external bus at the same time, snoop inquiries related to the second transaction will not be generated until the first transaction has been globally observed.

## SUMMARY

[0020] Embodiments of the present invention provide a method of processing a bus transaction in which an address is retrieved from the bus transaction and referred to a queue of pending transactions. A match indicator signal is returned from the queue. If the match indicator signal indicates a match, a snoop probe for the bus transaction is blocked.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0021] FIG. 1 is a block diagram of a conventional multi-agent system.

[0022] FIG. 2 is a block diagram of a bus sequencing unit of an agent constructed in accordance with an embodiment of the present invention.

[0023] FIG. 3 is a flow diagram illustrating operation of a snoop queue in accordance with an embodiment of the present invention.

[0024] FIG. 4 is a block diagram illustrating relevant portions of an external transaction queue and a snoop queue constructed in accordance with an embodiment of the present invention.

## DETAILED DESCRIPTION

[0025] The present invention alleviates the disadvantages of the prior art by providing an agent having a snoop queue whose depth is independent of the pipeline depth of its external bus. Embodiments of the present invention provide a snoop queue with a snoop blocking function that is coordinated with an external transaction queue. When the snoop queue observes an external bus transaction, before it issues a snoop probe for cache coherency checks, it refers the address of the new transaction to the external transaction queue. The external transaction queue compares the address of the new transaction with addresses of earlier-posted transactions that have not yet been globally observed. If a match occurs, the external transaction queue identifies the match to the snoop queue which in turn, blocks a snoop probe for the new transaction. After the pending transaction has been globally observed, the block is released.

[0026] In an embodiment, the principles of the present invention may be applied in a bus sequencing unit 200 ("BSU") of an agent, shown in FIG. 2. The BSU 200 includes an arbiter 210, an internal cache 220, an internal transaction queue 230, an external transaction queue 240 and the snoop queue 250. An external bus controller 300 interfaces the BSU 200 to the external bus 60. The BSU 200 fulfills data requests issued by, for example, an agent core 100.

[0027] The arbiter 210 receives data requests from not only the core 100 but also from a variety of other sources such as the snoop queue 250. Of the possibly several data requests received simultaneously by the arbiter 210, the arbiter 210 selects and outputs one of them to the remainder of the BSU 200.

[0028] The internal cache 220 stores data in several cache entries. It possesses logic responsive to a data request to

determine whether the cache 220 stores a valid copy of requested data and, if so, it furnishes the requested data in response thereto.

[0029] The internal transaction queue 230 receives and stores data requests issued by the arbiter 210. It coordinates with the internal cache 220 to determine if the requested data "hits" (was furnished by) the internal cache 220. If not, if a data request "misses" the internal cache 220, the internal transaction queue 230 forwards the data request to the external transaction queue 240.

[0030] The external transaction queue 240 interprets data requests and generates external bus transactions to fulfill them. The external transaction queue 240 is populated by several queue entries. The external transaction queue 240 manages the agent's transactions as they progress on the external bus 60. For example, when data is available in response to a transaction, the external transaction queue 240 retrieves the data and forwards it to, for example, the core 100.

[0031] The snoop queue 250 performs cache coherency checks within the agent. Typically, in response to a new bus transaction issued by another agent, the snoop queue 250 generates snoop probes to various caches within the agent (such as internal cache 220) and to the internal and external transaction queues 230, 240. It receives responses to the snoop probes and generates snoop responses therefrom. If necessary, the snoop queue 250 manages implicit writebacks of modified data from the agent.

[0032] The external bus controller 300 drives signals on the external bus as commanded by the external transaction queue 240 and snoop queue 250.

[0033] FIG. 3 illustrates a method 1000 of the snoop queue 250 operating in accordance with an embodiment of the present invention. It may begin when another agent requests data in a bus transaction. When a new transaction is posted, the snoop queue 250 decodes the transaction (Step 1010). It determines whether the transaction requires a cache coherency check. If so, the transaction requires a snoop probe (Step 1020). The snoop queue 250 then provides the address of the requested data to the external transaction queue 240 (Step 1030). Based upon a response from the external transaction queue, the snoop queue determines whether the address of the new transaction matches the address of a posted transaction (Step 1040). If so, the snoop queue blocks a snoop probe related to the new transaction (Step 1050).

[0034] Eventually, the prior conflicting transaction will be globally observed. When that occurs, the snoop queue releases the block (Step 1060). It emits a snoop probe within the agent and generates a snoop response according to conventional techniques (Step 1070).

[0035] If, at Step 1040, no match occurred, the snoop queue 250 advances to Step 1070 and emits the snoop probe.

[0036] FIG. 4 is a partial block diagram of the external transaction queue 240 and the snoop queue 250. The external transaction queue 240 is populated by a number of queue entries ("ETQ entries") 242. For each pending bus transaction posted by the external transaction queue 240, one of the ETQ entries 242 stores information regarding the transaction. Such information may include the request type, the

address of the transaction and/or the current phase of the transaction. The address field of each ETQ entry 242 includes match detection logic 244. The external transaction queue also includes observation logic 246 in communication with the match detection logic 244 and with the snoop queue 250.

[0037] During operation, the external transaction queue 240 receives an address of a new transaction from the snoop queue 250. The observation detection logic 246 forwards the received address to each match detection logic 244. It also observes outputs of the match detection logic 244 to determine whether the address stored in any ETQ entry 242 matches the received address. In the event of a match, the observation detection logic 246 reads the phase from the matching ETQ entry 242 and determines whether the matching transaction has already been issued onto the bus, but not yet been globally observed. If so, the observation detection logic 246 signals to the snoop queue that a conflict match exists.

[0038] The snoop queue 250 is also populated by a plurality of entries ("snoop queue entries") 252. The number of snoop queue entries 252 is independent of the pipeline depth of the external bus 60. It is also independent of the number of ETQ entries 242. The snoop queue 250 possesses control logic 254 to implement the method of FIG. 3. It forwards the address of new transactions to the external transaction queue 240. The control logic 254 also receives the match signal from the external transaction queue 240. Each snoop queue entry 252 includes a blocking bit (not shown) which, if enabled, prevents the snoop queue 240 from issuing a snoop probe. Responsive to a match signal from the external transaction queue, the control logic 254 enables the blocking bit. The blocking bit remains enabled until the pending conflicting transaction is globally observed. Thereafter, the bit is cleared and a snoop probe may be issued.

[0039] In an embodiment, each of the ETQ entries 242 is assigned a unique identifier ("ETQ ID"). When a conflict match exists, the observation detection logic 246 may provide the ETQ ID of the conflicting transaction to the snoop queue 250.

[0040] In an embodiment where the external transaction queue 240 furnishes the ETQ ID of a pending conflicting transaction, the snoop queue 240 may store the ETQ ID in a snoop queue entry 252 of the new transaction when it enables the blocking bit. In this embodiment, when the EBC 300 receives snoop responses, it forwards them to both the external transaction queue 240 and the snoop queue 250. The EBC 300 relates the snoop response to a transaction using its ETQ ID. Upon receipt of the snoop responses and the ETQ ID, the snoop queue 250 releases the blocking bit of all snoops which were being blocked by the associated ETQ transaction.

[0041] Optionally, the snoop queue 250 may be configured to ignore certain types of transactions. For example, a conflicting write back transaction does not raise coherency issues for a subsequent transaction because global observation of the write transaction does not necessarily mean that the agent is giving up ownership of the cache line. Also, an "uncacheable read," one that causes an agent to read but not cache requested data, does not cause state changes to occur within the agent when the read transaction is globally observed. In this embodiment, the observation detection

logic 246 also reads the request type out of the ETQ entry 242 of the matching pending transaction. Further, a "self snoop," another transaction identified by its request type, need not block a transaction. The observation logic 246, based on the request type, may not indicate "block" even though an address match occurred with an outstanding transaction.

[0042] Thus the present invention provides a snoop queue having a reduced queue size. The snoop queue of the present invention severs the relationship between the depth of the snoop queue and the pipeline depth of the external bus. The snoop queue of the present invention includes a snoop probe blocking feature to eliminate the boundary conditions that may exist when two agent issue transactions requesting the same data.

[0043] Several embodiments of the present invention are specifically illustrated and described herein. However, it will be appreciated that modifications and variations of the present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.

We claim:

1. A method of processing a bus transaction, comprising:
  - retrieving an address from the bus transaction,
  - referring the address to a queued list of pending transactions,
  - receiving a match indicator signal from the queue, and
  - if the match indicator signal indicates a match, blocking a snoop probe for the bus transaction.
2. The method of claim 1, further comprising emitting the snoop probe when the matching pending transaction is globally observed.
3. The method of claim 1, further comprising receiving a queue entry identifier in association with the match indicator signal.
4. The method of claim 3, further comprising:
  - receiving snoop results of the matching pending transaction identified by the queue entry identifier, and
  - emitting the snoop probe when the snoop results of the matching pending transaction is received.
5. A method of processing a bus transaction comprising:
  - retrieving an address from the bus transaction,
  - forwarding the address to a queue of pending transactions,
  - receiving a match indication signal and a request type signal from the queue,
  - based on the match indication signal and the request type signal, blocking a snoop probe for the bus transaction.
6. The method of claim 5, wherein the blocking occurs when the match indication signal indicates a match with a pending transaction.
7. The method of claim 5, wherein the blocking step does not occur when the match indication signal indicates a match with a pending transaction and the request type signal indicates that the matching pending transaction is a write transaction.
8. The method of claim 5, wherein the blocking step does not occur when the match indication signal indicates a match

with a pending transaction and the request type signal indicates that the matching pending transaction is an uncacheable read of data.

9. The method of claim 5, further comprising emitting the snoop probe when the matching pending transaction is globally observed.

10. The method of claim 5, further comprising receiving a queue entry identifier in association with the match indicator signal.

11. The method of claim 10, further comprising:

receiving snoop results of the matching pending transaction identified by the queue entry identifier, and

emitting the snoop probe when the snoop results of the matching pending transaction is received.

12. A bus sequencing unit of an agent, comprising:

an external transaction queue to process external transactions of the agent, the external transaction queue coupled to an agent bus output and populated by a plurality of transaction queue entries, and a snoop queue to process cache coherency operations of the agent, the snoop queue coupled to the agent bus output and populated by a plurality of snoop queue entries, the number of snoop queue entries being independent of the number of transaction queue entries.

13. The bus sequencing unit of claim 12, wherein the agent bus output is an external bus controller.

14. The bus sequencing unit of claim 12, wherein the transaction queue entries include an address field and match detection logic coupled to the address field and to an externally applied address input.

15. The bus sequencing unit of claim 12, wherein the snoop queue entries include an address field.

16. The bus sequencing unit of claim 12, wherein the snoop queue includes decoding logic.

17. A method of processing a bus transaction, comprising:

at a snoop queue:

buffering the bus transaction, retrieving an address from the bus transaction,

forwarding the address to an external transaction queue,

at an external transaction queue:

determining whether the address matches an address of a pending transaction,

returning a match indicator signal to the snoop queue representing whether the address matches an address of a pending transaction, and

at the snoop queue, blocking a snoop probe if the match indicator signal indicates a match.

18. The method of claim 17, wherein the blocking occurs when the match indication signal indicates a match with a pending transaction.

19. The method of claim 17, wherein the blocking step does not occur when the match indication signal indicates a match with a pending transaction and the request type signal indicates that the matching pending transaction is a write transaction.

20. The method of claim 17, wherein the blocking step does not occur when the match indication signal indicates a

match with a pending transaction and the request type signal indicates that the matching pending transaction is an uncacheable read of data.

21. The method of claim 17, further comprising emitting the snoop probe when the matching pending transaction is globally observed.

22. The method of claim 17, further comprising receiving a queue entry identifier in association with the match indicator signal.

23. The method of claim 22, further comprising:

receiving snoop results of the matching pending transaction identified by the queue entry identifier, and

emitting the snoop probe when the snoop results of the matching pending transaction is received.

\* \* \* \* \*